

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (MM-YYYY) 10-2004		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Reducing Aircraft Down for Lack of Parts with Sporadic Demand				5a. CONTRACT NUMBER 0979-15-01-01	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Bachman, Tovey C				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) LMI 2000 Corporate Ridge McLean, VA 22102-7805				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Logistics Agency, Defense Supply Center Richmond Mr. Clifford Wolfe DSCR-VEE, Bldg 33G 8000 Jefferson Davis Highway Richmond, VA 23297-5611				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT In the military aerospace environment, certain repair parts are infrequently demanded, but stocked because they are essential to maintaining a weapon system critical to the war-fighter. Because of their sporadic demand, it is difficult to decide when to buy these items and in what quantities. As systems become more reliable and failure rates decrease, the number of these infrequently demanded parts is likely to grow. Earlier studies found the Peak ordering policy the author invented significantly reduced wholesale wait-time and backorders. Rigorous new experiments confirm the benefits of the Peak policy, and show it can reduce retail wait-time and backorders as well. By considering the distribution of retail backorders over an aircraft squadron, we estimate the resulting reduction in the number of aircraft down for lack of parts. We also analyze the policy's near-term effect on inventory value and procurement workload After 5 years of development and review, the Peak policy is mature enough for implementation.					
15. SUBJECT TERMS Availability, backorders, infrequent demand, inventory; low demand; ordering; peak policy; procurement; repair parts; supply; supply chain; sporadic; system support					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unclassified Unlimited	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON Tovey C. Bachman
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (include area code) 703-917-7361

Reducing Aircraft Down for Lack of Parts with Sporadic Demand

Tovey C. Bachman, Ph.D.

LMI

2000 Corporate Ridge
McLean, Virginia 22102-7805

Phone: 703-917-7361

FAX: 703-917-7477

tbachman@lmi.org

72nd MORS Symposium

Working Group 19

24 June 2004

Abstract

In the military aerospace environment, certain repair parts are infrequently demanded, but stocked because they are essential to maintaining a weapon system critical to the war-fighter. Because of their sporadic demand, it is difficult to decide when to buy these items and in what quantities. As systems become more reliable and failure rates decrease, the number of these infrequently demanded parts is likely to grow. Earlier studies for the Defense Logistics Agency (DLA) and the Federal Aviation Administration (FAA)—organizations that manage parts inventories for repairing complex systems—found the Peak ordering policy the author invented significantly reduced wholesale wait-time and backorders. Rigorous new experiments confirm the benefits of the Peak policy, and show it can reduce retail wait-time and backorders as well. By considering the distribution of retail backorders, or “holes,” over an aircraft squadron, we estimate the resulting reduction in the number of aircraft down for lack of parts. We also analyze the policy’s near-term effect on inventory value and procurement workload, showing that the Peak policy can reduce both within a few years of policy initiation. After 5 years of development and review, the Peak policy is mature enough for implementation. A live test is underway, and broader implementation is under consideration.

Background

For frequently demanded parts, there is a well-developed theory and set of processes for ordering that balances the investment in inventory with customer service. Unfortunately, when parts experience only infrequent demand, that theory breaks down, and the established processes no longer work well.

An inventory management system for a single site typically manages each item using two control levels: an item’s reorder point (ROP), which determines when to order, and a requisitioning objective (RO), which determines how much to order. An order is placed when assets on-hand plus on-order decrease to or below the ROP, and the difference between the RO and current assets is the quantity ordered. The RO is usually the ROP plus a nominal order quantity, Q , often a Wilson lot size (“economic order quantity”). Thus Q is the quantity ordered if assets drop exactly to the ROP.

The ROP is an estimate of lead-time demand plus a safety level that protects against variability in lead-time demand.

Safety-level computations usually treat the number of demands in a lead-time as a random variable with a tractable theoretical probability distribution (e.g., Poisson, negative binomial, Laplace, or normal), estimate the mean and variance, and derive expressions for expected backorders and inventory cost as a function of the safety level. Mathematical optimization techniques are then used to set item safety levels to balance inventory investment with expected backorders, probability of a stockout, or system availability.

This approach to optimizing ordering policies for a single site began in the 1950s (Galliher et al., 1959), and has developed to include a great variety of policies (Silver, 1998). It also has been extended to optimize policies across a supply chain (Kruse, 1979), account for repair actions as well as ordering actions (Sherbrooke, 1992; Slay et al., 1996), and treat items that apply to diverse weapon systems with distinct availability goals (O'Malley, 1983). When there is sufficient demand data to characterize the lead-time demand distribution, but theoretical distributions do not fit well, non-parametric techniques, such as the bootstrap method (Fricker and Goodhart, 2000), may apply.

But what happens when items experience long and irregular periods of inactivity between demands (6 months to several years), what we call sporadic demand? For these items, the lead-time demand is usually zero. In a previous study of more than 300,000 sporadic-demand items, the author found that nearly 95 percent of lead-time intervals contained no demand. Forecasting lead-time demand for these items is extraordinarily difficult, as is forecasting demand variance. That is why a theoretical demand probability distribution is impractical—mean and variance cannot be estimated in any meaningful way. Use of empirical demand probabilities is possible, but for many sporadic-demand items, the data are too sparse to build a reasonable lead-time demand distribution. For example, if an item's only observed demands in the last 5 years comprise a demand for 8 units and another demand for 50 units, there is no reason to believe a demand for 20 units has a probability of zero. Modern enterprise resource planning (ERP) systems of the commercial sector focus on frequently demanded items; since commercial firms do not typically stock sporadic-demand items, ERPs offer no solution.

Inventory management specialists have long sought a successful approach to setting ROPs and ROs for sporadic-demand items. Usually heuristic policies are employed; and the military services use more-or-less arbitrary levels. Such policies fail to link inventory investment to service level, and generally do not work well. There are more sophisticated approaches. Croston showed that, when there is a constant probability of demand in a time interval, high fill rates can be obtained by basing reorder points on forecasts of both the time of next demand and demand size (Croston, 1972; for a recent survey of articles on policies based on statistical forecast-based methods, see Silver, 1998). Kruse divided an item's population into subsets by pooling items with similar lead-times, prices, and demand frequencies; thereby, obtaining enough demand data for empirical lead-time demand probabilities. Kruse assigned each item subset a common ROP based on a fill rate goal (LMI documented his method earlier; see Bachman and Bosma, 2003). Unfortunately, when demand is as irregular as it is for DLA-managed sporadic-demand items, none of these approaches have been shown to improve service levels (e.g., reduce customer wait-time) without significantly increasing inventory investment.

The scope of the challenge facing DLA—setting cost-effective ROPs for sporadic-demand items—is enormous. The agency manages nearly 1 million aviation stock numbers with sporadic demand, the inventory of those items is valued at more than \$1.5 billion and annual sales are in excess of \$400 million (Bachman and Bosma, 2003). Although no one item is typically active in any given year, sporadic-demand items experience significant aggregate activity and investment. Furthermore, the lack of DLA parts can render critical weapon systems inoperable—an event with significant military readiness consequences.

Over the last 5 years, the author developed the Peak ordering policy for sporadic-demand items, with the goal of reducing customer wait-times without increasing inventory investment. The most recent work, documented here, was accomplished with the capable support of a number of LMI employees.

Overview

The focus of this paper is on consumable parts, those deemed uneconomical to repair—when they fail on a weapon system, they are simply replaced. The supply chain for consumable parts is hierarchical, with the weapon system maintainer—the customer—at the end of the chain. To repair an aviation system, the maintainer requests parts from a local supply activity, what we call retail supply. (Retail supply belongs to the military service that owns and maintains the weapon system.) Retail supply requisitions parts from DLA, its wholesale supply organization. In turn, DLA buys parts from its vendors. The section “Aircraft-Level Analysis” illustrates this supply chain.

This paper describes the Peak ordering policy the author developed for DLA’s sporadic-demand items and its potential benefits at the wholesale level and the aircraft level, where the lack of a part may ground a weapon system. (While the focus is on aircraft items, there is nothing in this work that suggests the results apply only to aviation). The first two sections describe DLA’s current policy—the baseline—and the Peak policy. The paper then describes the simulation analyses we performed at the wholesale level to compare investment, customer service, and procurement workload for baseline and wholesale policies. The fourth section discusses the near-term effect of converting from DLA’s current practice to the Peak ordering policy. The paper then describes simulations of the Peak policy’s effect on part shortages at the aircraft level. The conclusion discusses the status of live testing.

Baseline and Peak Policies

Inventory management systems that distinguish between items with more regular demand and those with sporadic demand typically use a three-part policy. One part is an ordering policy for replenishment items (i.e., those replenished regularly) with statistical forecast-based ROPs and ROs. The second part is an activity threshold, typically set in terms of historical requisition frequency and quantity, which separates replenishment items from sporadic-demand items. Part three is a heuristic ordering policy employed for items with activity levels below the threshold.

DLA separates its replenishment items from its sporadic-demand items (which it refers to as numeric stockage objective [NSO] items), with an activity threshold that is based on last year’s demand. If an item experiences at least 4 requisitions and at least 12 units demanded in the trailing year, DLA uses a replenishment policy that treats lead-time demand with a Laplace distribution. The agency uses a smoothed forecast of lead-time demand to estimate the mean and a smoothed average of absolute forecast errors to estimate the variance. It also employs a modified Wilson order quantity, and

computes each item's expected backorders as a function of the item's order quantity and safety level. DLA then employs Lagrange multiplier optimization to set safety levels across items; thus minimizing the inventory cost required to meet an expected backorder constraint (Presutti and Trepp, 1970).

For NSO items (i.e., items with demand activity below the threshold), DLA uses a heuristic policy that sets the RO to the demand quantity in the trailing year and the ROP to half the RO. This policy does not link inventory cost to service level, however; and special rules apply to small subsets of items (Bachman and Bosma, 2003), which are not considered here. Items also may migrate between replenishment and NSO status quarterly.

The Peak policy changes the baseline policy in two ways:

- It introduces a new activity threshold between sporadic demand and replenishment, as well as a new way to apply the threshold. Setting this threshold properly is the key to improving service levels with the Peak policy while controlling inventory value (Bachman and Bosma, 2003).
- It introduces a new ordering policy for the sporadic-demand segment of the item population. In the context of the Peak policy, we no longer call these NSO items to emphasize the new segmentation.

The new threshold generally results in fewer replenishment items and more sporadic-demand items than the current threshold allows, but there is no change in the ordering policy for replenishment items.

The Peak policy utilizes quarterly demand data, as does DLA's current policy. We define an item's demand frequency in a time interval as the fraction of quarters with demand, irrespective of quantity. For example, if an item has 2 quarters with demand in a 5-quarter interval, the demand frequency is 0.4—the demand in one of those quarters may be for 1 unit, and in the other it may be for 100 units. The Peak policy relies on an activity history for each item, which is created from the quarterly demand history that DLA already maintains as follows: Each quarter's demand quantity is replaced by one when that quantity is positive; demand quantities of zero are unchanged. Each quarter, the policy applies a single exponential smoothing forecast (Brown, 1963; Sherbrooke, 1992) to this activity history in order to forecast future demand frequency.

When an item's forecasted demand frequency meets or exceeds the activity threshold, DLA's replenishment policy applies. If the frequency forecast is below the threshold, the new ordering policy applies. This decision is illustrated in Figure 1.

Using an item's quarterly history of demand quantity (DLA maintains at least 10 years' worth), we define its Peak demand as the maximum quarterly demand in a trailing K -year period. This period is called the look-back (see Figure 2). The look-back may vary by item population, but it is constant across items within a population.

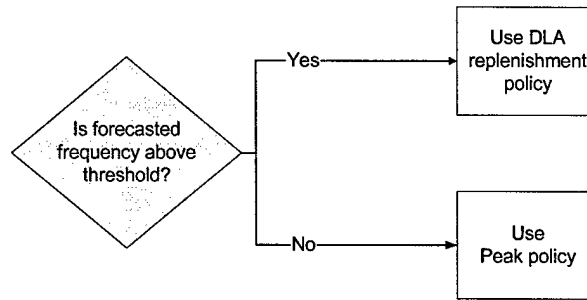


Figure 1. Using Frequency Threshold to Determine Ordering Policy

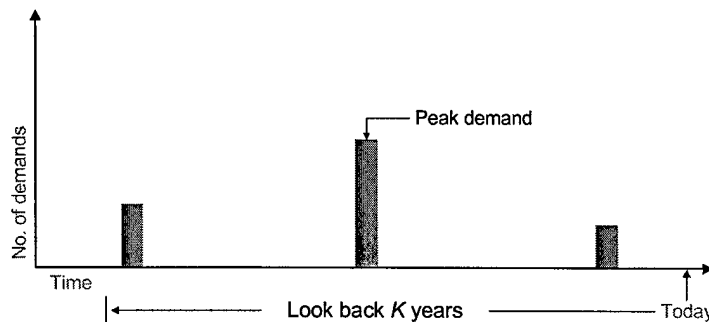


Figure 2. Computing Peak Demand

An item's ROP is the product of a price-based multiplier and its Peak demand, where higher multipliers apply to inexpensive items and lower multipliers apply to more expensive ones. For example, an item with a unit price of \$2.59 might use a multiplier of 5.0, but an item with a unit price of \$3,100 might use a multiplier of 0.3—in other words, we can afford a greater level of protection against backorders for the first item than we can for the second. The policy employs a set of price-based order quantities, in which cheaper items receive larger quantities than more expensive items, and sets an item's RO to its ROP plus its order quantity. Larger order quantities for inexpensive items allow us to avoid excessive administrative costs and workload that would result from frequent procurements. Assets on-hand and on-order are compared to the ROP in a continuous review. When assets are at or below the reorder point, Peak policy orders enough to bring assets up to the RO.

Peak Policy Development

With the possible exception of its frequency forecast, the Peak policy may appear simple; however, it is not at all trivial (and it was unclear at the policy's inception whether it was even possible) to find values of control parameters that achieve a high level of service without increasing inventory investment or procurement workload. This paper shows that one can make three-way tradeoffs between customer wait-time, inventory investment, and the number of procurement actions for a wide variety of sporadic-demand item populations. The process is iterative, and is made possible through the use of LMI's Financial and Inventory Simulation Model (FINISIM).

FINISIM provides three modes in which to generate item demand:

- A replay of demand history (retrospective simulation)
- Generation of synthetic demand patterns using empirical demand distributions
- A Poisson distribution.

It models the response of the inventory and related financial systems to demand patterns, emulating a wide variety of operating policies. In particular, FINISIM emulates DLA's current ordering policies, including item migration between NSO and replenishment classification over time and alternative policies tailored for sporadic demand. FINISIM's event processing algorithms are specifically engineered for rapid analyses when events are sparsely distributed over time. In such cases, FINISIM is as much as two orders of magnitude faster than other simulations. This speed is crucial in analyzing numerous alternative policies and finding values of control parameters that achieve a particular performance tradeoff. The model's customer service metrics include average customer wait-time, average backorder duration, number of backorder occurrences, average outstanding backorders, and both unit and requisition fill rates. Financial and workload metrics include annual dollars spent on procurements, number of procurement actions, and average value of on-hand inventory.

We use FINISIM's empirical distribution mode to develop Peak policies. In this mode the model generates long synthetic demand patterns (e.g., 200 quarters, or 50 years) in which the relative frequency of demands of different sizes, including zero, is close to the frequencies in the actual item demand history. For sporadic-demand items, these demand patterns are too sparse to build a realistic lead-time demand distribution; however, the demand patterns help adjust an ordering policy to respond to a demand pattern with a given frequency of zero demands and a given maximum demand.

There is one limitation of this approach, as far as projecting policy performance is concerned: Each item's generated demands form a stationary stochastic process. That is, there is the same set of probabilities for the number of units demanded at each point in time. A rigorous test of the Peak policy is discussed later in the section "Wholesale Analysis."

Using FINISIM to test numerous Peak policies on more than 10 item populations, each with 3,000–15,000 items, led to the conclusion that the best smoothing constant for the frequency forecast is 0.2, and the most cost-effective threshold is 0.6. Although the Peak policy was originally conceived to apply to items with gaps in demand of a year or more, this result shows the policy actually performs well on items with demand in as many as three out of five quarters. This result appears to be independent of the DLA item population; however, other control parameters must be tuned to the item population to achieve a given three-way performance tradeoff.

Although this process varies with the item population, it generally proceeds as follows:

- 1) Rank all items by unit price to determine the 25th, 50th, and 75th percentile prices.
- 2) Assign initial Peak multipliers of 1 for each of the resulting price quartiles, and set all order quantities to 1, as well. Call this policy Peak 1.
- 3) Use FINISIM to estimate the resulting on-hand inventory value, customer wait-time, and number of procurement actions for the item population.
- 4) Compare the performance of Peak 1 with FINISIM's assessment of baseline policy.
- 5) Stop here if the on-hand inventory value from Peak 1 is no more than that of the baseline, the customer wait-time is significantly lower than that of the baseline, and the average number of procurement actions per year is close to the baseline. This usually is not the case; more often, one of the three metrics is not within the desired range (the behavior is dependent upon a combination of item demand patterns and prices.)
- 6) If the wait-time reduction is not large enough, but the inventory value is lower than the baseline, use a peak multiplier of 2 for the bottom 50th price percentile and leave the top 50th price percentile with a multiplier of 1. If wait-time reduction is significant, but the on-hand inventory value is too high, try a multiplier of 1 for the bottom 50th price percentile and a multiplier of 0.5 for the top 50th price percentile. In either case, call this policy Peak 2.
- 7) Use FINISIM to compare Peak 2 results with baseline policy. If the metrics are in range, stop; if not, introduce new multipliers (e.g., 4, 2, 1, or 0.5) for all price quartiles.
- 8) Continue refining the values of multipliers in this way until there is a clear reduction in customer wait-time and the inventory value is less than that of the baseline.
- 9) With order quantities of 1, procurement actions often exceed the baseline. If so, introduce price-based order quantities, with larger quantities for the lower price quartiles (e.g., 20, 4, 2, or 1). Because this increases on-hand inventory value, reduce the Peak multipliers to counteract the effect.
- 10) Follow another iterative process to obtain order quantities (and reduced multipliers) that yield procurement actions that are no higher than the baseline, an on-hand inventory value that is no more than the baseline, and a significant reduction in customer wait-time.

To accommodate an especially high demand for the least expensive items, add a bottom 5th percentile price category; if investment is driven heavily by the items with the highest price, add a separate category for items above the 95th price percentile. Although the success of the above process has not been proven mathematically, it has worked for a wide variety of item populations. Peak policies that result from this process are very efficient in terms of customer wait-time per dollar value of inventory, but we cannot claim they are optimal.

Wholesale Analysis

In developing a Peak policy for an item population, FINISIM projects certain levels of customer wait-time, inventory value, and procurement actions. Because these projections are based on synthetic item demands (as described in the previous section), it was appropriate to construct a rigorous experiment to compare Peak and baseline policies by taking long item-demand histories, developing a Peak policy based on an initial segment of those histories, and assessing it based on the

remaining part of the histories. The demand data used to develop Peak policies would thus have different demand probabilities from the demands in the subsequent assessment period. If a policy did well, it would not be because of bias in the assessment method.

The author selected populations of DLA-managed items that apply to five critical weapon systems: the AH-64 Apache, E-2C Hawkeye, E-3 Sentry, C-5 Galaxy, and F/A-18 Hornet. We used 9-year quarterly item demand histories for these items, beginning with the first quarter of 1995. In some of the populations, aggregate demand increased over time; in others, it decreased. Items were limited to those with a unique application to each weapon system so it would be clear which weapon system would benefit from any improvement in supply performance. From each population, a sporadic demand subset was extracted, consisting of items that experienced demand in no more than 6 out of the earliest 10 quarters of their demand histories. Each sporadic-demand item population is referred to by the name of the associated weapon system. For example, "C-5 items" designates the sporadic-demand items that apply to the C-5.

For each of the item populations, we used the first 4 years of demand history and FINISIM to develop several Peak policies, each with a different set of objectives. One policy, "closest cost match," matched the dollar value of on-hand inventory with that of the baseline policy while decreasing customer wait-time, increasing fill rates, and keeping procurement actions no higher than the baseline. Another Peak policy, "relax orders constraint," had the same objectives, except it allowed for procurement actions to exceed the baseline. The policy "closest performance match" sought to keep customer wait-time and fill rates close to the baseline while reducing both inventory value and procurement actions. The policy "high performance" had the goal of reducing wait-time increasing fill rates, and allowed inventory value and procurement actions to increase in order to boost service levels. A "compromise" Peak policy sought to balance improvements in wait-time with reductions in inventory value and procurement actions.

We created a set of Peak policies that offered a three-way tradeoff between improved service levels (i.e., shorter wait-time and higher fill rates), reduced inventory value, and reduced procurement actions, based on the first 4 years of demand histories. Generally one or two metrics could be improved while constraining a third; however, not every policy option was available across all five populations. As always the ability to develop a particular policy depended upon the joint distribution of item prices and demands in a population.

To assess Peak policies, we performed retrospective simulations with FINISIM. Each assessment employed the baseline policy from the first quarter of 1995 through the last quarter of 1998—we made no policy change during the development period. We then continued in two ways:

- Start using a Peak policy at the end of the first 4 years, allow 2 years for simulated procurements to arrive, and then measure performance in the last 3 years.
- Continue using DLA's baseline policy to the end of the 9-year period, measuring performance in the last 3 years.

Figure 3 illustrates our experimental design.

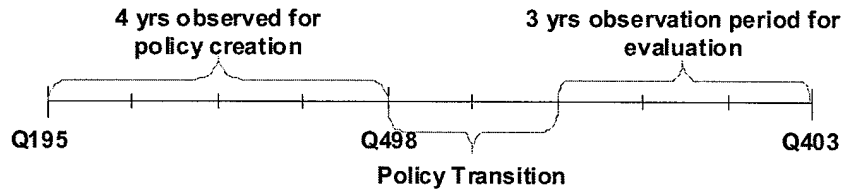


Figure 3. Experimental design

Figure 4 illustrates the results for Peak policies developed for E-2C items, which was a “middle-of-the-road” case in terms of improvement over the baseline. Each policy name (shown at the top of the chart) refers to the objective used to develop the Peak policy. The grouped bars show performance for each Peak policy according to four metrics. Each metric is the percentage difference between the Peak and baseline policies’ performance, based on averages over the 3-year assessment period. Lower numbers are better for the first three metrics: wait-time, number of orders placed, and value of on-hand inventory. Higher numbers are better for the fourth metric: unit fill rate. Baseline policy performance, in absolute terms, is shown to the right of each chart.

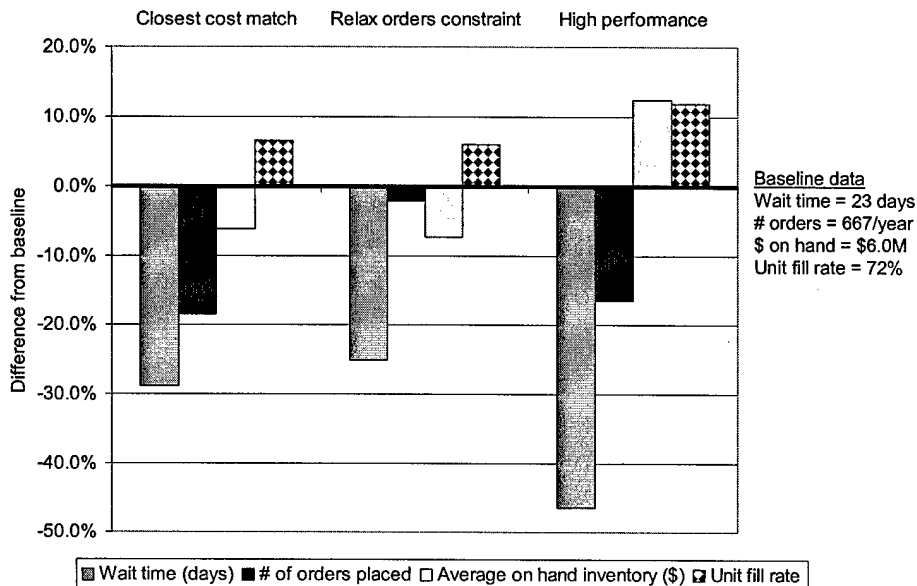


Figure 4. Assessment Results for Peak Policies Developed for E-2C Items

The first Peak policy, “closest cost match” produced a nearly 30 percent reduction in wait-time, and a nearly 20 percent reduction in procurement actions, relative to DLA’s current policy. Although it was developed to match inventory value with the baseline, the assessment demonstrated an inventory value reduction of approximately 5 percent, which was better than expected.

As expected, the reduction in procurement actions in the second option, “relax orders constraint,” was less than that of the first Peak policy. Allowing more procurement actions than the first Peak policy resulted in a slightly greater reduction in inventory value than the first policy. Wait-time reduction was still about 25 percent, but it was not as large as that of the first option. So, relaxing the constraint on the number of procurement actions did not produce a significant benefit in the other two metrics relative to the first Peak policy.

The “high performance” Peak policy, developed to increase performance significantly, achieved its objective—wait-time was reduced about 45 percent, the number of procurement actions declined, and the value of inventory increased by a little more than 10 percent. Our results indicated a range of Peak policies—each emphasizing different performance metrics—are available for the E-2C items.

Figures 5 through 8 illustrate the Peak policy assessment results for the other four weapon systems. The policies shown are only illustrative; many other policies could be developed.

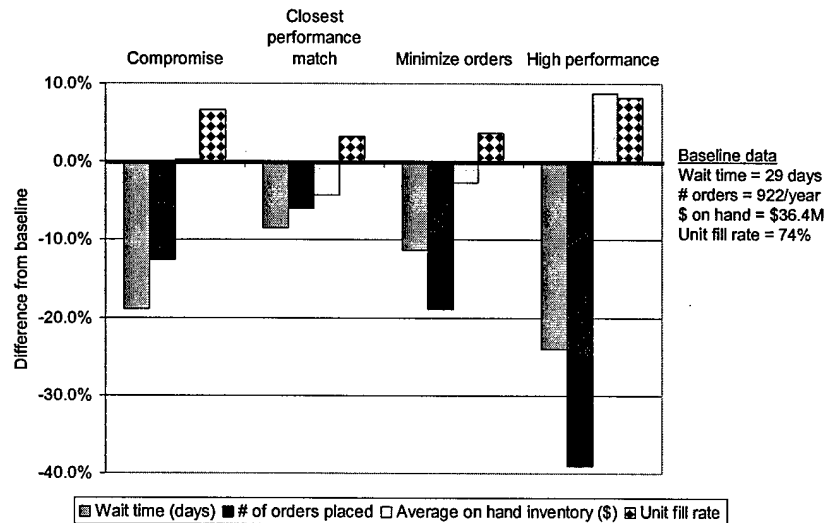


Figure 5. Assessment Results for Peak Policies Developed for AH-64 Items

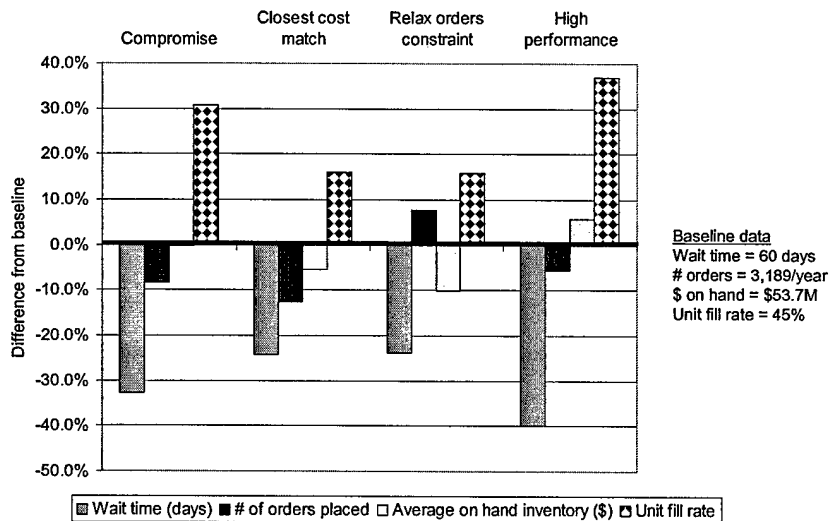


Figure 6. Assessment Results for Peak Policies Developed for C-5 Items

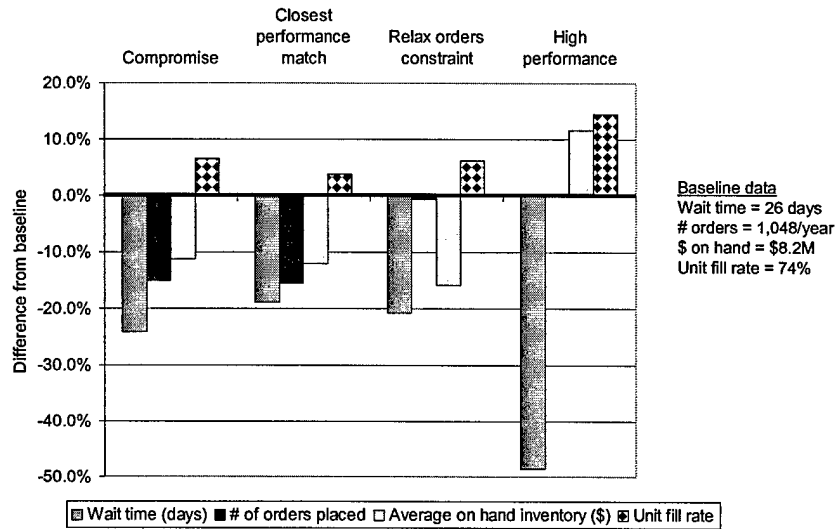


Figure 7. Assessment Results for Peak Policies Developed for E-3 Items

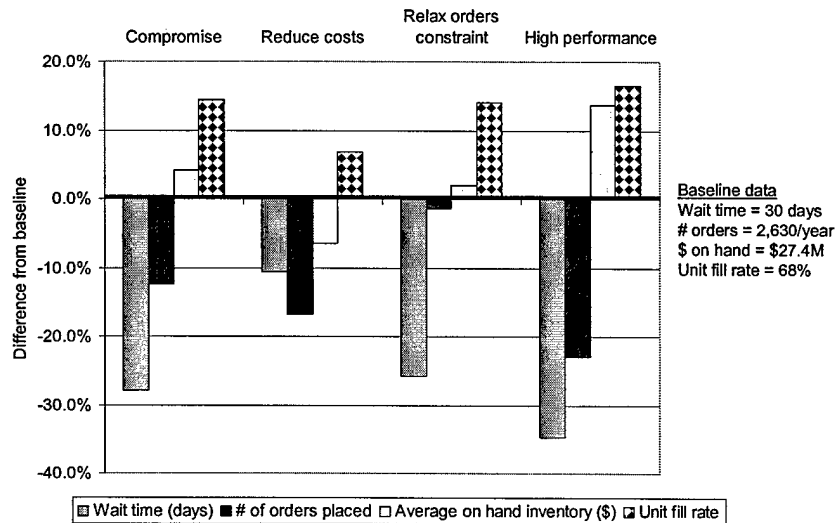


Figure 8. Assessment Results for Peak policies Developed for F/A-18 Items

Assessments of item populations generally demonstrate that policies behave in a manner consistent with their objectives.

- Peak policies developed to reduce wait-time did so in the assessments.
- Peak policies that sought to reduce procurement actions achieved that goal.
- Policies that attempted to reduce inventory value kept it at or below the value produced by the baseline policy.

For all item populations, we found policies that offered a nice compromise among the competing objectives: reducing wait-time, reducing the number of procurements, and reducing inventory value.

We performed separate analyses to test the threshold condition for item migration. The goal was to determine if there were any negative effects (for example, a large increase in inventory value) if we started with the full item populations (rather than populations with initial demand frequency that did not exceed 0.6) and let the Peak policy's threshold control item migration between replenishment and sporadic demand. These separate analyses followed the same timeline for policy conversion and assessment as those shown in Figure 3. We observed no negative effects from allowing the new threshold and the frequency forecast decide when Peak ordering should apply.

Taken together, these results confirm the Peak policy's benefits. And they show that, for the first time, DLA can make three-way performance tradeoffs for sporadic-demand items.

Near-Term Impact

Assessment results in the previous section were 3-year averages, taken over a period that starts after a Peak policy has been in effect for 2 years. It is natural to ask, "What happens to inventory value and procurement actions when policy conversion occurs?" Short-term increases in these two measures are inevitable when parts are purchased according to a new ordering policy. After all, DLA is purchasing a different mix of items than what the previous policy has "put on the shelf."

We used FINISIM to analyze "compromise" policies considered in the previous section, projecting how inventory value and procurement actions change over time. Figures 9 and 10 illustrate the on-hand inventory value and the number of procurement actions by year for the compromise policy developed for C-5 items.

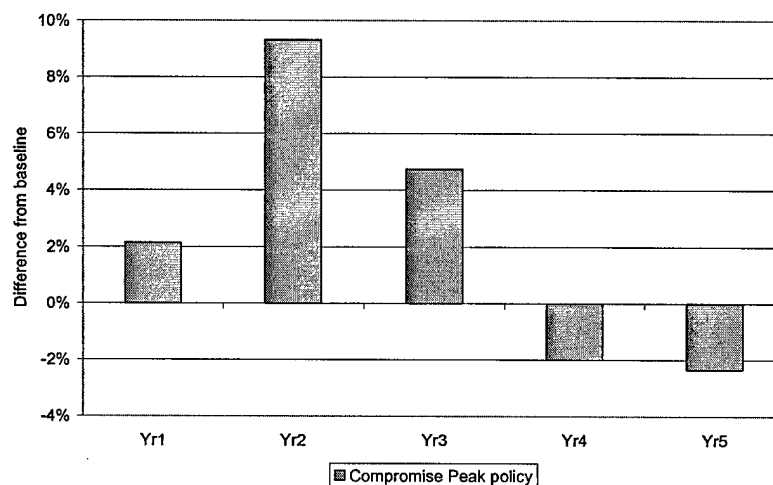


Figure 9. Value of On-Hand Inventory by Year for C-5 Compromise Peak Policy

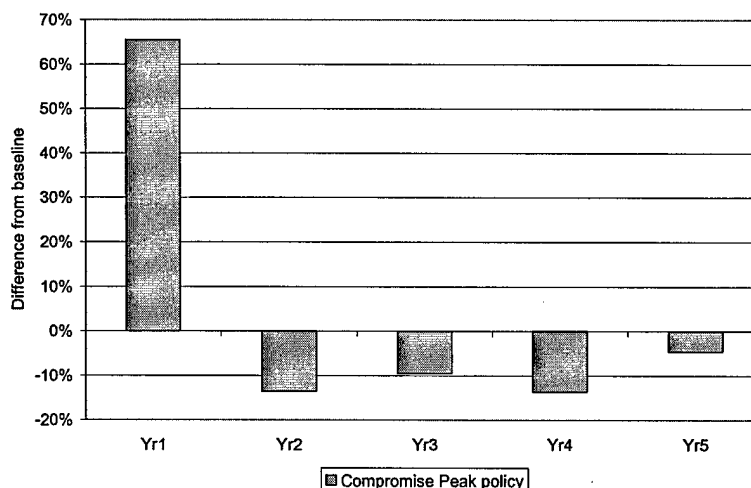


Figure 10. Number of Procurement Actions by Year for C-5 Compromise Peak Policy

As seen in Figure 9, on-hand inventory was up only slightly over the baseline in year 1. This reflects many procurement lead-times of more than a year. By year 2, inventory value was up significantly, which reflects the arrival of most material ordered at policy conversion. Moving from year 2 to year 3, the value of inventory declined—sales exceeded buys for the first time. This continued as we moved from year 3 to year 4, when inventory value fell below the baseline. Inventory remained below the baseline in year 5, which reflects the purchase of a better mix of spare parts—more of what is bought sells. The C-5 items served as a “middle case.” With some weapon systems it took longer for the inventory value to begin declining; with others it began sooner.

Figure 10 illustrates the expected surge of procurement actions in year 1, as the new policy triggered the buys necessary to produce a new mix of inventory. In year 2 and thereafter, annual procurements were below those of the baseline policy. So, the increase in procurement workload was one-time event, and was consistent across the item populations. Policies developed by LMI for DLA moderate the year-1 procurement workload and spending.

Aircraft-Level Analysis

The wholesale analysis complete, we turned to examining the effect of changing the wholesale ordering policy on parts shortages at the end of the supply chain—the aircraft. An aircraft rendered inoperable (“down”) for lack of a part is said to be not mission capable due to supply (NMCS); and the occurrence of a backorder, or “hole,” at the aircraft level is referred to as an NMCS incident. We proved that using the Peak policy at the wholesale level could reduce NMCS incidents.

Because personnel at the aircraft level often resolve an NMCS incident by cannibalizing a part from other weapon systems or finding other workarounds, we could not claim an increase in actual readiness; however, it is DLA’s job to make the required parts available, not to assume its customers will work around parts shortages. For this reason, it was appropriate to measure the effectiveness of the Peak policy through the increase in aircraft availability (the percentage of the fleet not down for lack of a part), ignoring workarounds. We called this “no-cannibalization availability,” although other workarounds are discounted as well.

We used populations of DLA items that apply to three types of aircraft: the Navy's E-2C and F/A-18, and the Air Force's E-3. For the first two aircraft, we used items the Navy had identified as first indentures (items removed directly from the aircraft, as opposed to items removed from a reparable component). First indenture items served as proxies for items that could ground the aircraft. For the E-3, we used items that grounded the aircraft at some point in the past. This analysis was not limited to items that had a unique application to the subject aircraft—we included items common to multiple weapon systems. Item populations, and the Peak policies developed for them, differ from those discussed earlier in "Wholesale Analysis," in which the goal was to show we could produce a three-way performance tradeoff.

From these three item populations, we built wholesale demand histories and extracted sporadic-demand items (i.e., items with demand in no more than 6 out of the first 10 quarters). We then developed Peak policies for each item population using the first 5 years of quarterly item demand histories (from the beginning of 1995 through end of 1999). The goal of these policies was to significantly reduce wholesale customer wait-time, but keep inventory value and procurement actions near the baseline.

Using a retrospective simulation, we replayed demands for the first 5 years, employing the baseline policy. We activated the Peak policy at the start of 2000, and allowed it to run for 2 years so that most of the assets from simulated procurements could arrive. We passed each item's simulated wholesale assets (on-hand, due-in from procurement, and backorders) and wholesale levels (ROP and RO) as of the end of 2001 to a multi-echelon supply chain simulation (a version of FINISIM that projects backorders at the aircraft level). Another set of final wholesale assets and levels, which were produced using the baseline policy from the beginning of 1995 to the end of 2001, was also passed to the multi-echelon FINISIM for comparison. Figure 11 illustrates the experimental design.

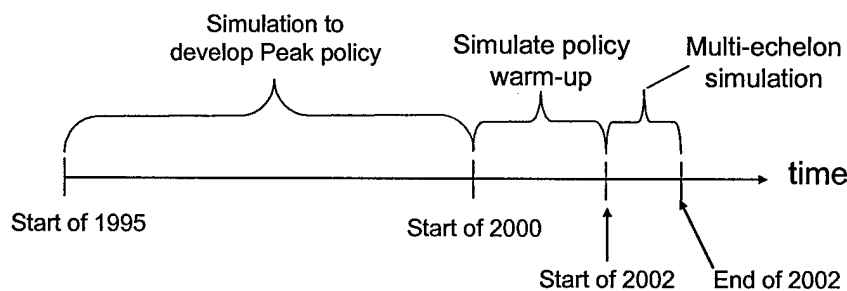


Figure 11. Experimental Design for Aircraft-Level Analysis

To examine the effect of Peak policies on DLA's Navy customers, we chose four aircraft carriers, two from the Atlantic Fleet and two from the Pacific Fleet. The carriers and their associated aircraft deckloads (for 2002) were

- USS *George Washington* (4 E-2Cs, 36 F/A-18s),
- USS *John F. Kennedy* (4 E-2Cs, 24 F/A-18s),
- USS *Kitty Hawk* (4 E-2Cs, 36 F/A-18s), and
- USS *John C. Stennis* (4 E-2Cs, 36 F/A-18s).

For Air Force customers, we considered the three permanent Air Force bases for the E-3:

- Elmendorff AFB
- Kadena AFB
- Tinker AFB.

Of these, only Tinker AFB had significant demand for DLA-managed sporadic-demand items in 2002, so we only performed a multi-echelon simulation for that base. Tinker was assigned 28 E-3 aircraft in 2002.

We analyzed backorders at the aircraft using a daily, multi-echelon simulation for 2002. The lowest echelon of the supply chain was the ultimate customer for parts in our item populations—an aircraft fleet (i.e., one or more squadrons). The next echelon up was the local supply activity, or retail supply. Above that, we modeled wholesale supply, with the top echelon DLA’s vendors (i.e., suppliers), which we treated as a single entity. Figure 12 illustrates the simulation for the Navy case.

Simulations for the other aircraft-ship combinations were similar. For the E-3, Tinker AFB replaces the aircraft carrier in Figure 12. At the retail level (Tinker AFB), other customers include demands from depot activities as well as other aircraft based at Tinker AFB. At the DLA level, other DLA customers represent wholesale demands from all locations other than Tinker.

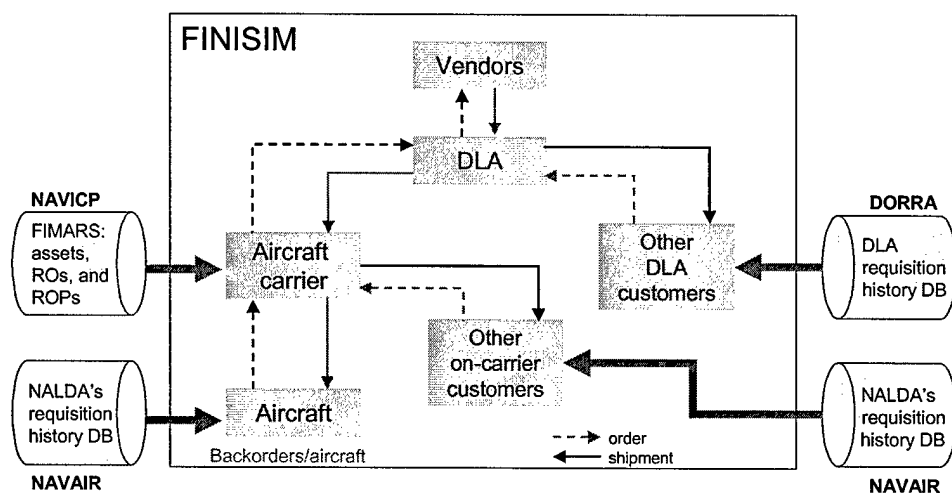


Figure 12. Structure of Multi-Echelon Simulation for Navy Customers

Consider the case of the E-2C squadron on the *John C. Stennis*. We used Navy maintenance data for parts required to replace items removed directly from the E-2C to model requests on the ship's supply activity (from aircraft to aircraft carrier in Figure 12). We also used Navy maintenance data to model demands from other repair activities competing for the same parts (from other on-carrier customers to aircraft carrier in Figure 12), to repair another aircraft type on the same carrier, for

example. We emulated the Navy's ordering policy using their levels data; so, when assets were low enough, FINISIM generated a simulated order from the carrier to DLA. We modeled competing wholesale demand, from customers other than the *Stennis*, using DLA requisition history data (from the other DLA customers to DLA in Figure 12.) FINISIM then generated simulated demands for parts from DLA to its vendors. We measured performance at the aircraft, aircraft carrier and DLA levels.

We assessed the effects of Peak policies on the aircraft in two ways:

- We projected the average number of aircraft down for lack of parts at a location by simulating the average number of outstanding backorders at the aircraft level, and considered the resulting holes as uniformly distributed over the aircraft at that location. Any aircraft with at least one hole was down for a part (it was NMCS).
- We used our no-cannibalization availability formula (shown in Equation 1) to estimate the probability that a randomly chosen aircraft was available at the location.

$$A_0 = \prod_{i=1}^N \left(1 - \frac{BO_i}{NAC} \right)$$

where A_0 is availability, the percentage of the fleet not down for a part, (Equation 1.)

N is the number of parts, BO_i is the number of backorders for part i ,

and NAC is the number of aircraft.

Each factor in the product is our estimated probability that an aircraft is not lacking a particular part. Treating holes for different parts independently, the product of all such factors is the probability of an aircraft not being down for any of the subject parts.

For the F/A-18s on two carriers, changing from the baseline to Peak policy at the wholesale level reduced part shortages and increased no-cannibalization aircraft availability. For the *George Washington*, we reduced average outstanding backorders from 6.35 to 5.36. Distributing the resulting holes uniformly over the 36 aircraft, we had an average of 6 aircraft down with DLA's baseline policy; we had an average of 5 aircraft down with the Peak policy, and the no-cannibalization availability increased by 3 percent (i.e., the difference between availability with the baseline policy and availability with the Peak policy, where each is computed using Equation 1). On the *Kennedy*, the Peak policy reduced the average outstanding backorders from 3.50 to 1.62—two fewer aircraft down due to the lack of a part—and increased availability by 7 percent.

We observed no change in backorders at the aircraft level for F/A-18s on the other two carriers or E-2Cs on any carrier (no positive or negative effects). This was expected; after all, worldwide demand for the parts in question is sparse, even over several years. This means that, at any given customer location, only a few parts are active in a single year. Although the Peak policy makes parts more available for a population of items, taken as a whole, it does not improve performance for every item—it may not help the few active parts at one customer location. Even if the Peak policy increases availability for certain parts at the wholesale level, there is no guarantee any particular customer will benefit in a 1-year period—other customers may get the parts first, which leaves no parts for the customer analyzed.

For the Air Force case, the Peak policy reduced average outstanding E-3 backorders from 29.73 to 27.40—a reduction of two holes. (This excludes one part with extraordinarily large demand in 2002; with this part included, average outstanding backorders would have been reduced from 104 to 85). Using the baseline policy, we had 30 holes on average. Distributing these holes uniformly, all 28 aircraft were down (26 aircraft have 1 hole and 2 aircraft have 2 holes). With the Peak policy, the average number of holes was 27.4, so we can anticipate either one less aircraft down or no improvement, depending on whether the 27.4 represents 27 holes or 28 holes. In either case we reduced part shortages at the aircraft-level (as we saw in two of the F/A-18 experiments) and no-cannibalization aircraft availability, as defined by Equation 1, increased by 8.6 percent.

From this analysis, we concluded the Peak policy decreased the number of aircraft down for lack of sporadic-demand DLA parts. In two cases it was a decrease of one or two aircraft; and in a third case, it was one less aircraft down, but only for about half the time. In no case did the Peak policy increase the number of aircraft down or parts shortages.

We believe the benefit of the Peak policies is understated, because we did not capture the effect of making DLA-managed parts more available for reparable item repair, which could reduce holes for those items, as well.

Our analysis is significant in another way: To the best of our knowledge, these were the first experiments to model the effects of DLA wholesale supply policy on supply-oriented readiness measures using realistic customer demand. Modeling demand properly is critical to accurately projecting performance. In previous LMI experiments, wholesale backorders can be too low (by as much as a factor of three) if theoretical demand probabilities, rather than empirical demand data, are used.

Live Testing

In February of 2004, the DLA's Defense Supply Center, Richmond (DSCR), asked LMI to develop a Peak policy to improve support for 16,000 sporadic-demand items that apply to 15 key aircraft and engines. A joint effort involving DLA and the military services selected specific items for various fixed-wing aircraft (A-10, E-2, E-3, EA-6B, C-5, F-15, F/A-18, and S-3), helicopters (AH-64, CH-47, HH-60G, and UH-60), and engines (F100 series, F404, and TF-39).

The Peak policy we developed is projected to reduce customer wait-time by 35 percent by year 2 and by 60 percent in the longer term, all while staying within DSCR's guidelines for initial spending and number of procurement actions. Working with DSCR, we produced a variant of the policy that could be implemented within DLA's ordering system, with similar performance and cost parameters.

Starting in March 2004, DSCR began to implement this revised Peak policy, electing to test it on C-5 items. Because many items have procurement lead-times in the range of 1 to 3 years, we expect results of this test to begin emerging by late summer of 2005. Headquarters DLA is also reviewing the Peak policy for possibly wider implementation.

Conclusion

We developed the Peak ordering policy for items with sporadic-demand patterns, and a simulation model that enables us to make three-way tradeoffs between the resulting service level, value of inventory, and procurement workload. The ability to make tradeoffs—long available for more frequently demanded items—is new for sporadic-demand items. We showed this capability enables us to produce Peak policies that reduce wholesale customer wait-time by 20 to 45 percent while maintaining or reducing wholesale inventory value and procurement actions. We further demonstrated that the Peak policy can reduce parts shortages at the weapon system level and reduce the number of aircraft down for lack of a part.

After 5 years of development and review, the Peak policy is now mature enough for implementation. A live test is underway, and broader implementation is under consideration.

Acknowledgements

The author wishes to thank DLA's Aging Aircraft Program Office for supporting this research; DLA's Office of Operations Research and Resource Analysis (DORRA) for providing wholesale supply data; the Naval Air Systems Command (NAVAIR) for providing Navy maintenance data; the Naval Inventory Control Point, Philadelphia, for Navy retail supply data; and the Air Force Logistics Management Agency (AFLMA) for Air Force retail supply and maintenance data. The author also wishes to thank Carol DeZwarte of LMI for checking the data and performing many of the simulation analyses. Production of this manuscript was partially supported by the LMI Research Institute.

References

- Bachman, T.C., and Bosma, L.F., *Defense Logistics Agency's Numeric Stockage Objective Policy* (Report SVD10T4). McLean VA: LMI, 2003.
- Brown, R.G., *Smoothing Forecasting and Prediction*. Englewood Cliffs, NJ: Prentice Hall, 1963.
- Croston, J.D., "Forecasting Stock Control for Intermittent Demands," *Operations Research Quarterly*, vol. 23, no. 3 (1972): 289–303.
- Fricker, R.D., and Goodhart, C.A., "Applying a Bootstrap Approach for Setting Reorder Points in Military Supply Systems," *Naval Research Logistics*, vol. 47, no. 6 (2000): 459–478.
- Gallagher, H.P., Morse, P.M., et al., "Dynamics of Two Classes of Continuous Review Inventory Systems," *Operations Research*, vol. 7, no. 3 (1959).
- Kruse, W.K., *An Exact N-Echelon Inventory Model: The Simple Simon Method* (Technical Report 79-2). U.S. Army Inventory Research Office, 1979.
- O'Malley, T.J., *The Aircraft Availability Model: Conceptual Framework and Mathematics*. McLean, VA: LMI, 1983.
- Presutti, V.J., and Trepp, R.C., "Much Ado About Economic Order Quantities (EOQ)," *Naval Research Logistics Quarterly*, vol. 17, no. 2 (1970).
- Sherbrooke, C.C., *Optimal Inventory Modeling of Systems: Multi-echelon Techniques*. New York: Wiley, 1992.

Silver, E.A., Pyke, D.F., et al, 1998. *Inventory Management and Production Planning and Scheduling*, Wiley, New York.

Slay, F.M., Bachman, T.C., et al., *Optimizing Spares Support: The Aircraft Sustainability Model*, (Report AF501MR1). McLean, VA: LMI, 1996.